REMARKS

Applicants appreciate the thorough examination of the present application that is reflected in the Official Action of August 18, 2004. Applicants also appreciate the Examiner's citation of U.S. Patent 6,516,337 to Tripp et al. Applicants have now carefully studied Tripp et al., and respectfully submit that all of the pending claims are patentable over Tripp et al. for the reasons that now will be described.

The Hyperlinks Have Been Removed

In response to the objection to the specification, the hyperlinks have been changed by removing "http:\\www" or "www" from the hyperlinks, so that there no longer are executable hyperlinks in the specification. The drawings have not been changed to eliminate hyperlinks, because the drawings are only available as image files, so that executable hyperlinks are not found in the drawings. Accordingly, the objection to the specification over the use of hyperlinks has been overcome.

Moreover, the specification has been reviewed and amended for proper usage of trademarks as appropriate. Accordingly, this objection to the specification has been overcome.

Claim Amendments

The claims have been amended to correct informalities, as will now be described. In particular, Claims 1, 20, 21, 25, 26 and 32 have been amended to eliminate recitations of "steps" and to provide proper antecedent basis. Claims 34 and 35 have been amended to eliminate "means for" language in these computer program product claims.

Claims 1-14 and 16-31 Are Patentable Over Tripp et al.

Independent Claim 1 stands rejected under 35 USC §102(e) over Tripp et al. Claim 1 recites:

> 1.   A method of efficiently serving content in a distributed computing environment, comprising:
>   receiving usage metrics for a particular stored object; and
>   evaluating the received usage metrics to determine whether the particular stored object is stored in an appropriate location, and moving the particular stored location if not.

Accordingly, content may be efficiently served in a distributed computing environment by determining whether a particular stored object is stored in an appropriate location based on usage metrics for the particular stored objects. The particular stored location may then be moved, i.e., the object may be moved to another stored location, if the particular stored object is not stored in an appropriate location, based on the usage metrics that are received. The location of content storage may thereby be determined intelligently based on usage metrics for a particular stored object.

In sharp contrast, Tripp et al. relates *"Sending to a Central Indexing Site Meta Data or Signatures From Objects on a Computer Network"*, as noted in the Tripp et al. title. As noted in the Tripp et al. Abstract:

> A search engine utilizes a bottom-up approach to index the content of a network instead of relying on atop-down [sic] approach as used by conventional search engines. The network being indexed may be any network, including the global computer network, Internet, and the world wide web. Instead of using a central site including spidering software to recursively search all linked web pages and generate a search index of the Internet, independent distributed components or agents are located at each web site and report meta data about objects at the web site to the central server. A central catalog of object references is compiled on the central site from the meta data reported from each web site. One or more brochures file [sic] may also be created and stored on each web site to provide conceptual or non-keyword data about the site, such as demographics and categorization information. This conceptual information is then utilized in constructing the central catalog so that more accurate search results may be generated for search queries applied to the catalog.

Thus, Tripp et al. relates to indexing of a web site, to allow web searching using search engines. Rather than using spiders to search the web, independent distributed components or agents are located in each web site and report metadata about objects at the website to a central server. Accordingly, Tripp et al. does not appear to be related to solving the problem of determining whether a particular stored object is stored in an appropriate location, as recited in Claim 1, but, rather, appears to be directed to techniques for indexing of web sites.

As noted in the above-quoted section of the Abstract of Tripp et al., a brochure is used for indexing. As noted at Column 11, lines 15-34 of Tripp et al.:

> The brochure **206** is a small file that may contain conceptual and other non-HTML information which would be useful to improve the indexing of sites or parts of a single site on the remote server **208**. A brochure **206** may contain any information pertinent to the web site,

including but not limited to keywords, phrases, categorizations of content, purpose of the site, and other information not generally stored in a web page. The brochure **206** is generated manually by individual web site administrators. The administrator fills out a form at the central server **202**, and receives an email containing the brochure **206** or downloads the brochure after submitting the form contents. Upon receiving the brochure **206**, the administrator stores it within the file structure of the web site on the remote server **208**. There may be multiple brochures **206** at the same web site, each describing specific portions of the site. Each brochure **206** may refer to a single web page or a group of web pages stored within a specific subdirectory at the web site. All information stored in each brochure **206** is applied to the pages referenced in the brochure.

Moreover, as noted in Column 20, lines 26-32:

The agent **204** generates a site index, which is a database. The database includes a number of tables, each table consisting of records (rows) and fields (columns). Each table in the database includes similar records to speed searches. All Tables may be sorted alphabetically and then by category. In one embodiment of the agent **204**, the agent generates Tables 3-12 as shown below.

One of the key words shown in Tables 3 and 4 is a "ranking". As noted in Tripp et al.

Column 30, lines 40-62:

The indexing system **200** may also perform ranking of web pages having references in the central index. First, the agent **204** may perform positional and contextual rankings for particular words in the web pages on a site. The positional rankings assign a ranking value to a word based upon, for example, the location of the word in the web page and the position of the word relative to other words in the page. The contextual ranking is determined using contextual information about the site contained in the brochure **206**. For example, if a word in a web page corresponds to a category as listed in the brochure **206** (see Table 2), the word will be assigned a higher ranking. In addition to rankings generated by the agent **204**, the central server **202** also generates rankings for the central index. For example, the central server **202** may generate rankings based upon whether a page is a source or reference to the desired data. Rankings may also be determined based upon user input such as the usage or popularity of a site as measured by how often the site is linked as the source site in other sites, or through positive comments entered by users about the context or ranking of a site. All the methods of ranking just described are know [sic] as static rankings, meaning that the ranking is determined before a particular search query is applied.

Assume, for the sake of argument, that the above passage of Tripp et al. suggests some form of "receiving usage metrics for a particular stored object", as recited in Claim 1. However, there does not appear to be any description or suggestion in Tripp et al. of:

> evaluating the received usage metrics to determine whether the particular stored object is stored in an appropriate location, and moving the particular stored location if not,

as recited in Claim 1. This lack of description or suggestion is consistent with the fact that Tripp et al. is directed to a bottoms-up indexing approach for a network, and does not appear to be concerned with a determination of whether the particular stored object is stored in a particular location. For at least this reason, Claim 1 is patentable over Tripp et al.

In alleging that Tripp et al. describes or suggests the above-quoted recitation of Claim 1, the Official Action cited Column 16, lines 9-32 of Tripp et al., at Page 3, line 3 of the Detailed Action. However, this passage of Tripp et al. states:

> If the host or web site URL is not currently being indexed, the web server performs the following operations. First, an automatic email is sent to contacts at the host to encourage the host to install the agent. An automatic email is also sent to a contact person for the web site with a "Thank You" and a request that they ask their host to install the agent. In addition, a retrieval order is generated for the central server to retrieve the brochure file from the web site in one hour. If the retrieval order is unsuccessful, it will be repeated 2, 4, 8, 24 and 48 hours later, until successful. If still unsuccessful after 48 hours, the retrieval order is canceled. By verifying the presence of the site brochure in the specified location, unauthorized information about a site may not be created by a third party in an attempt to have their site indexed along with a more popular site. This is a common problem with existing search engines where a third party copies the keywords from a meta tag in a popular site. The bogus site with copied keywords is then submitted to a search engine for indexing, and when search queries are applied to the search engine that produce the popular site the bogus site is also produced. This may not be done with the site brochure because the brochure is not an html page available to outside persons and because it is encrypted so even if the file is obtained the information contained therein is not accessible.

With all due respect, this passage relates to how indexing is performed at the web server, by asking a web site to index itself. There is no description or suggestion of determining whether a particular stored object is stored in an appropriate location, or of determining whether the particular stored object is stored in an appropriate location by evaluating the received usage metrics, as recited in Claim 1. Finally, there is no description or suggestion of moving the particular stored location if the particular stored object is not stored in an appropriate location as recited in Claim 1. Accordingly, the recitations of Claim 1 are not described or suggested by Tripp et al.

Claims 2-13 and 16-31 are patentable at least per the patentability of Claim 1 from which they depend. In view of the clear patentability of Claim 1, a detailed analysis of these dependent claims will not be performed.
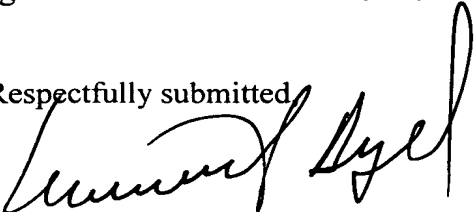
Claims 15 and 32-35 Are Patentable

Claims 15 and 32-35 are patentable at least for the same reasons that were described above in connection with Claim 1. Moreover, the Applicant-applied prior art that was cited in rejecting Claims 32-35 does not supply the missing teaching in Tripp et al. of:

> evaluating the received usage metrics to determine whether the particular stored object is stored in an appropriate location, and moving the particular stored location if not,

as recited in Claim 1.

Conclusion

Applicants again appreciate the Examiner's thorough analysis and the citation of Tripp et al. Applicants have now shown, however, that Tripp et al.'s bottom-up approach to index the content of a network by a search engine does not describe or suggest evaluating received usage metrics to determine whether a particular stored object is stored in an appropriate location, and moving the particular stored location, if not, as recited in the independent claims. Accordingly, all of the pending claims are now in condition for allowance, which is respectfully requested.

Respectfully submitted,

Mitchell S. Bigel
Registration No. 29,614

USPTO Customer No. 46589
Myers Bigel Sibley & Sajovec, P.A.
P.O. Box 37428
Raleigh, NC 27627
Telephone: (919) 854-1400
Facsimile: (919) 854-1401